

Communication Protocol for Passive Optical Network

Topologies

Inventors:

Onn Haran

Ariel Maislos

Lior Khermosh

RELATED APPLICATION

This application claims priority benefit of
provisional U.S. Patent Application Ser. No. 60/272,967,
filed 2 March 2001, entitled Ethernet PON; and priority
benefit of provisional U.S. Patent Application Ser.
No. 60/297,171, entitled Ethernet PON, which applications
are hereby incorporated by reference for all purposes.

BACKGROUND OF THE INVENTION

1. Field of the Invention

The present invention relates generally to point-to-
multipoint telecommunication systems. More particularly,
the invention relates to packet-based protocols for bi-
directional data interchange on a Passive Optical Network
(PON).

2. Background

With the explosion of the Internet content and other multimedia services, telephone line modems are becoming progressively less adequate as means for connecting
5 personal computers to high-speed networks. Several access systems with higher bandwidth capabilities have evolved, including Integrated Services Digital Network (ISDN), Digital Subscriber Line (DSL), and cable modems. These and similar systems are typically bandwidth-limited to
10 rates of the order of several megabits per second. It is desirable to increase the connection rates further to satisfy the growing end-user demand for bandwidth.

Because of its many advantages, optical fiber is widely used in telecommunication systems. The advantages
15 include low signal attenuation, immunity to electromagnetic interference, low crosstalk, fast propagation speed, physical flexibility, small size, low weight, and, most important, high bandwidth. Bringing optical fiber to the end-users is therefore desirable for
20 a number of reasons, including the higher bandwidth that optical fiber would make available to the end-users. But the expense of installing optical fiber in place of the

copper wires that typically connect each individual end-user to the central office remains high.

Passive optical network topology provides a good match with the need to connect multiple users to a central point. A passive optical network is an optical transmission system that brings optical fiber to the end-user without requiring individual fiber optic links from each end-user to the central office. Depending on its points of termination, a passive optical network can be called fiber-to-the-curb (FTTC), fiber-to-the-building (FTTB), or fiber-to-the-home (FTTH) network. A passive optical network typically includes a central controller, such as an optical line terminal, at the communication company's office, and a plurality of remote network nodes, such as optical network units or customer premises equipment, located near end-users of the network. In a PON, the central controller and the remote network nodes are interconnected by optical fiber and optical splitters/combiners ("splitters" hereinafter).

Figure 1 illustrates the topology of a representative passive optical network 100. Reference numeral 110 designates an optical line terminal (OLT) that communicates with multiple optical network units (ONUs)

over the PON 100. The signals broadcast from the OLT 110 travel through optical fiber spans 140 and passive splitters 120, and are received by the ONUs 130. Each ONU 130-N corresponds to a discrete end-user or application.

- 5 The signals sent by the ONUs 130 travel in the reverse direction to the OLT 110.

Note that the PON topology differs from the "star" configuration where each end-user is connected to a central location by a dedicated circuit, which may be a
10 dedicated physical or virtual circuit. The topological differences are important for at least two reasons. First, the OLT 110 broadcasts (point-to-multipoint) over PON 100, so that every ONU 130-N receives the broadcasts. Broadcasting raises privacy concerns because each ONU
15 receives not only the transmissions intended for it, but also all other transmissions from the OLT 110. Second, the uplink transmissions (from the ONUs 130-N to the OLT 110) must be multiplexed because the multiple ONUs 130-N must share the total bandwidth available on the shared
20 transmission medium of network span 140-1. Preferably, the available bandwidth should be allocated dynamically and adaptively.

Because of the need to multiplex and the need to secure downlink communications, protocols used in star topologies may not work well in PON topologies. A need therefore exists for a flexible and efficient communication protocol that will provide multiplexing for secure point-to-multipoint communications over a passive optical network topology.

SUMMARY OF THE INVENTION

10 In accordance with the principles of this invention, a protocol is provided for transferring data between a central controller and a first remote network node of a plurality of remote network nodes connected by a network with passive optical network topology. The remote
15 controller discovers the first remote node by sending a GATE control message to all remote nodes that have not been discovered. The GATE message includes a time stamp and a definition of a time period within which the undiscovered remote nodes may respond to the GATE message.
20 The first remote node receives the GATE message, sets its internal clock to the time stamp, and responds with a REGISTER_REQ control message that identifies the first remote node to the central controller. The central

controller receives the REGISTER_REQ message and sends a REGISTER control message to the first remote node. The first remote node thus becomes discovered. Thereafter, the central controller sends additional GATE messages to the first remote node, authorizing uplink transmissions from the first remote node at specific times. To preserve privacy of the downlink transmissions, the transmissions to the first remote node are encrypted with a key known to the first remote node, but is not known to other remote nodes residing on the network.

BRIEF DESCRIPTION OF THE DRAWINGS

The present invention will be explained, by way of examples only, with reference to the following description, appended claims, and accompanying figures where:

Figure 1 illustrates the passive optical network topology;

Figure 2 illustrates the process of discovery of a remote network node by a central controller;

Figure 3 illustrates back-to-back uplink data transmissions from two different optical network units;

Figure 4 illustrates a diagram of an OLT's state machine controlling the processes of ONU discovery and detection of loss of connection;

Figure 5 illustrates a diagram of an ONU's state machine controlling the processes of ONU discovery and detection of loss of connection;

Figure 6 illustrates the process of encryption of a downlink packet comprising multiple N-byte payload blocks and a short payload block;

Figure 7 illustrates encryption key replacement process with a NEW_KEY message retransmission.

DETAILED DESCRIPTION

Referring again to Figure 1, the OLT 110 transmits packets or frames downlink to the ONUs 130-N. The downlink transmissions pass through the optical fiber spans 140 and passive splitters 120, reaching all the ONUs 130-N. Uplink packet or frame transmissions pass through only those splitters and spans that lie between the OLT 110 and the particular ONU that sent the particular transmission. For example, a transmission from ONU 130-4 would travel to the OLT 110 through spans 140-11, 140-6, and 140-1, as well as through optical splitters 120-2 and

120-1. Because of the physical properties of passive optical splitters, the OLT 110 receives the signal attenuated only by the small losses in the fiber and the splitters; other ONUs 130-N receive only faint reflections of the signal.

Downlink and uplink transmissions can be separated in various ways, for example, by using different wavelengths or separating the signals using the directional quality of light. Thus, the uplink and downlink transmissions may overlap each other.

To prevent collisions between uplink transmissions sent from different ONUs 130-N, a time-division multiplexing (TDM) scheme is used, with the different ONUs 130-N transmitting packets at different times. It follows that the internal clocks of the OLT 110 and the ONUs 130-N should be synchronized. Preferably, the internal clocks of the ONUs 130-N are synchronized to the internal real-time clock of the OLT 110. In a first exemplary non-limiting embodiment, the real-time clock has a period of 16 nanoseconds and is stored in a 32-bit real-time clock register. When the clock reaches its maximum count (e.g., "FFFF FFFF"), the clock wraps around and starts counting from all "0."

An ONU's clock is generally synchronized with the OLT's real-time clock when the ONU is "discovered" by the OLT. Often, the discovery and synchronization processes occur when the ONU comes on-line. In operation, the
5 sequence of discovery and clock synchronization proceeds as described below with respect to a second exemplary non-limiting embodiment.

When the ONU comes on-line, it is initially in the undiscovered state. In this state, the ONU may not
10 transmit packets to the OLT because any unauthorized transmission can collide with authorized transmissions from other, discovered ONUs. While in the undiscovered state, the ONU listens for a GATE message addressed to all undiscovered ONUs. The GATE message authorizes all
15 undiscovered ONUs to transmit to the OLT, in order to become discovered. A representative GATE message is graphically illustrated in Table 1 below. Note that the GATE message and other control messages can be characterized by their Media Access Control (MAC) opcodes.

TABLE 1 - graphical illustration of a GATE message

Field name	Length	Description
Source Address	6 bytes	Address of the OLT
Destination address	6 bytes	Multicast / Address of the ONU
Packet type	2 bytes	MAC control 0x8808
MAC control opcode	2 bytes	GATE_OPCODE
Timestamp	4 bytes	Value of current local clock
Number of grants	1 byte	Describes the number of grants included in this message (1 to 6). When the MSB is 1, then discovery indication is set
Grant 0 start time	4 bytes	The time that ONU can start transmit at
Grant 0 length	2 bytes	Maximal transmission time
...		
Grant N start time	4 bytes	Same description as Grant 0 start time
Grant N length	2 bytes	Same description as Grant 0 length

The GATE message includes the current OLT time stamp,
 5 i.e., the value of the OLT's real-time clock when the GATE message was sent. (We will refer to the OLT clock as the PON clock.) The GATE message further includes the grant start time and grant length values in the appropriate fields. The grant start time is the time (according to
 10 the PON clock) when the undiscovered ONUs are allowed to begin uplink transmissions, and the grant length is the length of time the undiscovered ONUs are allowed to transmit. As illustrated in Table 1, multiple

authorizations to transmit may be included in the same
GATE message.

The GATE message optionally includes OLT's
parameters, for example, the time required for the OLT to
5 lock onto a signal; OLT's automatic gain control circuit
adjustment time; and protocol constants, such as the
timers described below.

When the undiscovered ONU receives the GATE message,
it loads its internal clock with the received value of the
10 PON clock. From this point on, the ONU is synchronized
with the OLT. For improved clock synchronization, the ONU
may load its internal clock with the values of PON clock
whenever the ONU receives a time-stamped message from the
OLT.

15 The undiscovered ONU is allowed to transmit its
REGISTER_REQ message when the PON clock reaches the grant
start value. Typically, a REGISTER_REQ message includes
the ONU's time stamp and various negotiation parameters.
The negotiation parameters may include, for example, the
20 time required for the ONU's laser to stabilize after turn-
on; the time required to turn off the ONU's laser; and
protocol buffer sizes for uplink transmissions. The
REGISTER_REQ message may further include an

acknowledgement of the OLT's capabilities described in the GATE message. A representative REGISTER_REQ message is graphically illustrated in Table 2 below.

Table 2 - graphical illustration of REGISTER_REQ message

Field name	Length	Description
Source Address	6 bytes	Address of the ONU
Destination address	6 bytes	Address of the OLT / multicast address
Packet type	2 bytes	MAC control 0x8808
MAC control opcode	2 bytes	REGISTER_REQ_OPCODE
Timestamp	4 bytes	Value of current local clock
Capabilities	N bytes	Optional

For collision avoidance, the ONU may delay its transmission of the REGISTER_REQ message by a period in the range of 0 to ($\text{<grant length>} - \text{<REGISTER_REQ transmission time>}$). In the second exemplary non-limiting embodiment, the delay period is randomly and uniformly distributed within this range. As persons skilled in the art will recognize, the upper limit of the delay range guarantees that the transmission of the REGISTER_REQ message will not extend beyond the time interval allocated for uplink transmissions from the undiscovered ONUs.

If a collision between the REGISTER_REQ messages from multiple undiscovered ONUs does occur, the undiscovered ONUs back off and wait until the next regular GATE message is sent to them. The undiscovered ONUs then participate
5 in the discovery process from the beginning. The collision may be detected by the OLT if the OLT has collision detection hardware, and the OLT may increase the grant length value in the subsequent GATE message, thereby increasing the discovery window and decreasing the
10 probability of collisions. Alternatively, the collision can be ignored.

The OLT detects the existence of a previously undiscovered ONU when the OLT receives a valid REGISTER_REQ message. The identity of the ONU is known
15 from the source address field of the REGISTER_REQ message. In a third exemplary non-limiting embodiment, the OLT subtracts the value received in the time stamp field of the REGISTER_REQ message from the OLT's current real-time clock value to calculate the round trip delay between the
20 OLT and the ONU. The calculated delay corresponds to the distance between the two nodes.

In the final step of the discovery process, the OLT transmits a REGISTER message to the ONU. The REGISTER message directs the ONU to switch its state to discovered state, and may also include an acknowledgement of some of the negotiation parameters transmitted by the ONU in its REGISTER_REQ message. A representative REGISTER message is graphically illustrated in Table 3 below.

TABLE 3 - graphical illustration of a REGISTER message

Field name	Length	Description
Source Address	6 bytes	Address of the OLT
Destination address	6 bytes	Address of the ONU
Packet type	2 bytes	MAC control 0x8808
MAC control opcode	2 bytes	REGISTER_OPCODE
Timestamp	4 bytes	Value of current local clock
Capabilities	N bytes	Optional

When the ONU receives the REGISTER message, it switches its state to "discovered," as directed by the REGISTER message. The discovery process, illustrated in Figure 2, is now complete.

Generally, the GATE messages are used not only in the course of the discovery process, but also to allow the OLT to control uplink transmissions. In a fourth exemplary non-limiting embodiment, the OLT controls uplink data transmissions by performing scheduling calculation and

transmitting GATE messages toward the ONUs. An ONU is allowed to transmit only if it has been authorized by a grant received in a GATE message from the OLT. Collision-free transmissions from ONUs are guaranteed because the OLT assigns grants without overlap. When in discovered state, the ONU preferably starts to transmit immediately after its local clock equals to the grant start time.

This is illustrated in Figure 3, where ONU #1 transmits during time interval 310, and ONU #2 transmits during time interval 330. Note that the OLT inserts a guard band 320 between consecutive transmissions from different ONUs to prevent collisions that might otherwise result because of the small differences between the local clocks of the two ONUs.

As illustrated in Table 2 above, several grants can be packed into a single GATE message to increase bandwidth utilization. Each grant contains a grant start time value and a grant length value in appropriate fields, thereby informing the addressed ONU when to start its transmission and the maximum length of the transmission.

Each ONU may participate in the process of bandwidth allocation by requesting bandwidth in a REPORT message.

Note that the REPORT message is sent when the ONU receives

a grant, i.e., an authorization to transmit. The REPORT message may contain the number of bytes the ONU requests to transmit per local queue, and one or more priorities, such as the priorities defined in the IEEE 802.1 standard.

- 5 A representative REPORT message is graphically illustrated in Table 4 below.

Table 4 - graphical illustration of a REPORT message

Field name	Length	Description
Source Address	6 bytes	Address of the ONU
Destination address	6 bytes	Address of the OLT / Multicast
Packet type	2 bytes	MAC control 0x8808
MAC control opcode	2 bytes	REPORT_OPCODE
Timestamp	4 bytes	Value of current local clock
Queue #0 request	4 bytes	Number of bytes requested for priority queue #0
Queue #1 request	4 bytes	Number of bytes requested for priority queue #1
Queue #2 request	4 bytes	Number of bytes requested for priority queue #2
Queue #3 request	4 bytes	Number of bytes requested for priority queue #3
Queue #4 request	4 bytes	Number of bytes requested for priority queue #4
Queue #5 request	4 bytes	Number of bytes requested for priority queue #5
Queue #6 request	4 bytes	Number of bytes requested for priority queue #6
Queue #7 request	4 bytes	Number of bytes requested for priority queue #7

The OLT need not explicitly acknowledge the REPORT message. A smart scheduler in the OLT simply uses the reported values to optimize bandwidth allocation among the different ONUs, and the OLT sends GATE messages in
5 accordance with the scheduler's output. The scheduler need not consider any specific reported values, and may in fact ignore all the reported values by allocating bandwidth based, for example, on a predefined fairness algorithm, or on the end-users' service classifications.

10 In a fifth exemplary non-limiting embodiment, both the ONU and the OLT detect loss of physical or logical connection between them. To detect loss of connection, the OLT runs a timeout routine with timers corresponding to the individual discovered ONUs. The timeout routine
15 monitors receipt of REPORT messages (or of other messages) from the individual ONUs. After a message is successfully received from an ONU, the timeout routine resets the corresponding timer. If one of the timeout timers reaches a preprogrammed threshold value before being reset, the
20 OLT assumes that the connection with the ONU corresponding to the timer has been interrupted. The OLT then switches the ONU to undiscovered state.

The OLT is allowed to switch an ONU to the undiscovered state by issuing, for example, a GATE message to the ONU with a discovery indication directing the ONU to change its state. (A special control message may also
5 be used for this purpose.) When the ONU receives this message, it releases all its pending grants and switches its state to undiscovered. The ONU can then participate in the next discovery process.

The ONUs run their own timeout routines that monitor
10 the receipt of GATE messages. In a sixth exemplary non-limiting embodiment, an ONU resets its local timeout timer after each successfully received GATE message. If the local timeout timer reaches a preprogrammed threshold before being reset, the ONU assumes that the physical or
15 logical connection with the OLT has been lost. The ONU then changes its state to undiscovered and awaits the next GATE message sent as part of the discovery process.

Figures 4 and 5 illustrate the states and transitions of the OLT's and ONU's state machines that govern the

20 discovery and detection of connection loss processes.

In a seventh exemplary non-limiting embodiment, the data exchange between the OLT and the ONUs conforms to the Ethernet standard IEEE 802.3. The OLT encrypts the

downlink traffic to allow only the addressed ONU to understand the data packets sent. The downlink Ethernet packets are encrypted with a block encryption algorithm at the physical sublayer level. The algorithm conforms to

5 the Advanced Encryption Standard (AES) promulgated by the National Institute of Standards and Technology. Non AES-conforming encryption algorithms can also be used.

The block encryption of the seventh embodiment is described with reference to Figure 6. Each packet 610 is

10 broken into N-byte-long blocks 612-1 through 612-J, and the last block having n bytes ($n \leq N$). Each N-byte-long block is encrypted independently by encryption process 620, resulting in encrypted blocks 612-1' through 612-J'.

If the last block is a short block, i.e., $n < N$, then the

15 encrypted next-to-last block 612-J is encrypted again using encryption process 630, which may be the same as the encryption process 620. The short block is then XOR-ed with the twice-encrypted next-to-last block. In Figure 6, the XOR process is designated with numeral 640. The

20 resulting encrypted packet 610' comprises the encrypted blocks 612-1' through 612-J', and the encrypted short block 614'.

In the seventh embodiment, each ONU periodically generates new encryption keys and transmits them to the OLT along with the keys' identifiers (IDs), which may include the keys' sequence numbers. Every downlink packet transmitted by the OLT contains a header with the packet's encryption key ID and key sequence number. When the ONU receives a packet, it compares the received packet's key ID to the ONU's own key ID. If the key IDs match, the ONU decrypts the packet. Otherwise, the ONU discards the packet. In this way, privacy of communications is maintained despite the fact that the downlink transmissions are received by all ONUs.

The ONU initiates a new key replacement sequence by transmitting a NEW_KEY message to the OLT. The message contains the next key for the OLT to use in encrypting the messages intended for the ONU, and an identifier of the key, which can be the key's sequence number. The sequence number can be simply the current key's sequence number incremented by one. A representative NEW_KEY message is graphically illustrated in Table 5 below.

TABLE 5 - graphical illustration of a NEW_KEY message

Field name	Length	Description
Source Address	6 bytes	Address of the ONU
Destination address	6 bytes	Address of the OLT / Multicast
Packet type	2 bytes	MAC control 0x8808
MAC control opcode	2 bytes	NEW_KEY_OPCODE
Timestamp	4 bytes	Value of current local clock
New Key	16 bytes	New key to use
Sequence number	1 bytes	Sequence number of new key

Alternatively, the OLT can initiate a key replacement sequence by requesting a new key from the ONU in a

5 NEW_KEY_REQUEST control message.

The OLT does not explicitly acknowledge the NEW_KEY message. Instead, the OLT begins to use the new key to encrypt the downlink traffic, and specifies the new key sequence number (or other ID) in the headers of the

10 encrypted packets. The ONU recognizes the new key from the key's ID and begins to decrypt the traffic using the new key. If the OLT has not switched to the new key within a predefined time, the ONU retransmits the NEW_KEY message. In the seventh embodiment, the ONU periodically

15 replaces the encryption key.

The flow of the key replacement process with a NEW_KEY message retransmission is illustrated in Figure 7.

We have described the invention and some of its features in considerable detail for illustration purposes. Neither the specific embodiments of the invention as a whole nor those of its features limit the general principles underlying the invention. In particular, the invention is not limited to an optical network, to connecting ONUs or customer premises equipment to a network, or to the use of AES-conforming encryption algorithms. The specific control messages illustrated in the tables of this specification contain fields that are not necessary to the operation of the invention, and do not contain other fields that can be added to these control messages. Moreover, the lengths of the specific fields within the control messages can vary from the lengths shown in the Tables above. For example, the New Key field in the NEW_KEY message can be increased to accommodate longer encryption keys. Many additional modifications are intended in the foregoing disclosure, and it will be appreciated by those of ordinary skill in the art that in some instances some features of the invention will be employed in the absence of a corresponding use of other features. The illustrative examples therefore do not define the metes and bounds of

the invention, which function has been reserved for the following claims and their equivalents when considered in conjunction with the rest of this specification.